Fake It to Make It: Exploratory Prototyping in HRI

J.D. Zamfirescu-Pereira
zamfi@berkeley.edu
UC Berkeley
Berkeley, CA

David Sirkin
sirkin@stanford.edu
Stanford University
Stanford, CA, USA

David Goedicke
dg536@cornell.edu
Cornell Tech
New York, NY

RAY LC
ray.lc@cityu.edu.hk
City University of Hong Kong
Hong Kong

Natalie Friedman
nvf4@cornell.edu
Cornell Tech
New York, NY

Ilan Mandel
im334@cornell.edu
Cornell Tech
New York, NY

Nikolas Martelaro
nikmart@cmu.edu
Carnegie Mellon University
Pittsburgh, PA

Wendy Ju
wendyju@cornell.edu
Cornell Tech
New York City, NY

ABSTRACT

Exploratory prototyping techniques are critical to devising new robot forms, actions, and behaviors, and to eliciting human responses to designed interactive features, early in the design process. In this opinion piece, we establish the contribution of exploratory prototyping to the field of human-robot interaction, arguing research engaged in design exploration—rather than controlled experimentation—*should be* focused on flexibility rather than specificity, possibility rather than replicability, and design insights as incubated subjectively through the designer rather than dispassionately proven by statistical analysis. We draw on literature in HCI for examples of published design explorations in academic venues, and to suggest how analogous contributions can be valued and evaluated by the HRI community. Lastly, we present and examine case studies of three design methods we have used in our own design work: physical prototyping with human-in-the-loop control, video prototyping, and virtual simulations.

CCS CONCEPTS

• **Human-centered computing** → **Interaction design process and methods.**

KEYWORDS

design, prototyping, exploration, experimentation, evaluation, wizard of oz, virtual simulation, video prototyping, HCI, HRI

ACM Reference Format:

J.D. Zamfirescu-Pereira, David Sirkin, David Goedicke, RAY LC, Natalie Friedman, Ilan Mandel, Nikolas Martelaro, and Wendy Ju. 2021. Fake It to Make It: Exploratory Prototyping in HRI. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21 Companion)*, March 8–11, 2021, Boulder, CO, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3434074.3446909>

This work is licensed under a Creative Commons Attribution International 4.0 License.

HRI '21 Companion, March 8–11, 2021, Boulder, CO, USA.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8290-8/21/03. <https://doi.org/10.1145/3434074.3446909>

Figure 1: What is the right way to interact with a chair robot? Let us find the ways.

1 INTRODUCTION

The need for and use of low-cost prototyping and elicitation techniques in design is well-established in human-computer interaction [8, 16]. In human-robot interaction, however, discussion of design techniques such as physical, video, and virtual prototyping, including Wizard of Oz approaches, has inappropriately focused on meeting standards of verifiability and replicability [21, 35, 43, 58, 61] more properly applied to controlled experimentation in research than to research through design. These standards are useful when applied to controlled studies, but when applied to design explorations, they inhibit the broader dissemination of subjective design insights that have a unique value in their own right.

Exploratory prototyping can help researchers in “making the right thing”, where controlled experiments help researchers in “making the thing right” [9, 64]. When used for design exploration, the emphasis of rapid prototyping methods *should be* on flexibility, possibility, and design insights as incubated subjectively through the designer, rather than replicability, specificity, and dispassionately proven statistical results. Design insights come from the exploration of wide

design spaces, and from the rapid prototyping of instances within those design spaces—not from an exacting understanding of tightly-controlled instances. Disseminating these insights, together with the context in which they are generated, helps other designers and researchers use those insights to inform new designs, taking into account how their own contexts might differ from the original.

In this opinion piece [59], we start by arguing (§2) for the importance of exploratory prototyping techniques in devising new robot forms, actions, and behaviors, and in eliciting human responses to designed interactive features early in the design process. These methods are ideally low-cost, iterative, and should evolve to rapidly respond to the designer’s subjective sense of what works and what does not. Sharing the results of these explorations, even when they are not conclusive, can provide substantial scaffolding within a new design space for other researchers to stand on.

Assessment of design explorations using the standards of controlled research experiments prevents academic HRI researchers from sharing methods or insights early, when they are the most valuable, before hardware decisions have been finalized and are expensive to revisit. In describing the process used to generate design insights, researchers should not focus on *replicability of outcomes*, but rather on an *understanding of context* of the design process—context provides designers an understanding of when and how others’ insights can inform their own designs. Similar conversations have taken place in the broader HCI community, and we draw on those conversations—as well as examples (§3.1) of exploratory prototyping work in HCI and HRI—to offer alternative ways of assessing design explorations in the context of HRI (§3.2).

We then zoom in to present and examine case studies of three design methods: physical prototyping with human-in-the-loop control (§4), video prototyping (§5), and virtual simulations (§6). In each case, we have used these methods to author interactions for users to have with chair robots (“ChairBots”) that we are in the process of designing (Fig. 1). We share these case studies and other examples to illustrate how designers can use these kinds of activities for HRI research, and to help the broader HRI community understand what contributions from such work could and should be (§7).

2 EXPLORATORY PROTOTYPING

Exploratory prototyping, in seeking to find “the right thing” [9], requires researchers to explore wide design spaces, and to rapidly prototype within them—questioning assumptions about materials, users, and contexts. Buxton describes lightweight, inexpensive prototyping late in the design process—often by way of testing with potential users, to get the “design right”—in contrast to exploration of alternatives, early in the process, that encourages divergent ideation to get to the “right design” [9].

In the following discussion, we describe why we believe exploratory prototyping techniques are powerful, why they should be used by human-robot interaction researchers to generate early design insight, and why these insights should be considered by the broader HRI community—with an understanding of the ways these insights differ from those offered by traditional controlled experiments.

2.1 Figuring out the “right design”

The British Design Council’s double-diamond design process [5] sets out four phases in the design process: discover, define, develop, and determine. **Design exploration** is the primary goal of the third phase, where designers seek to maximize the design alternatives

presented to address the problems and opportunities found in the *discover* and *define* phases of the process. As depicted in Fig. 2, the development phase is a divergent process, intended to generate a wide set of alternatives and solutions. If the search space for possible solutions or solution features is not cast widely enough, the outcomes will be stunted or sub-optimal [9, 17]. The goal of prototyping methods is to help designers to generate more alternatives, more divergence in alternatives, and to enable them to elicit stronger and more detailed feedback from users [55]—critical in the case of social robots, but perhaps less so for the factory floor or industrial variety.

While in traditional product design the emphasis is often placed on physical form, the advent of computers and robots has broadened the space of exploration to include motion, action, and behavioral response. Hoffman and Ju [22] have described numerous methods for early design of robot movement, useful as the movements are necessarily constrained or enabled by the physical form of the machine.

2.2 Prototyping in HRI

The human-robot interaction community understandably views actual autonomous interaction with actual robots in actual space as the ideal, most informative manner of prototyping interaction.

However, interaction with actual robots in actual space costs a great deal of actual time and actual money, not to mention investment in design details as well as modes and contexts of interaction that may well be off-target. Moreover, putting the actual robots performing actual interaction in front of the actual users that will eventually be using the robot may be highly impractical, depending on the users or uses of the robot. It is often necessary and even preferable to use lower-cost, faster, and less actualized prototypes to develop robots and determine their behaviors, actions, and responses. The oft-heard critique that design explorations did not “use a real robot” belies the principle, well understood in traditional design and engineering contexts, that increased fidelity early in the design process does not make for better outcomes [20]. The value of low-cost and rapid prototyping is clear, but *low-fidelity* prototyping is less well known in HRI to be a critical practice used by designers to construct design knowledge, communicate, and explore ideas in order to make decisions later in the design process [45, 57].

Prototyping without an actual robot also helps robotics researchers understand where the lower bound lies: what feature set is *truly needed* in a robot? “Does this arm need three joints?” “Does this

Figure 2: The British Design Council’s double-diamond design process [5]

delivery robot need omniwheels?” “Does this turntable need high-power stepper motors?” We are liable to make use of these features if they are available to us, even when a lesser alternative would suffice.

Realizing that a “real robot” is not necessary to make key discoveries and contributions to what kinds of interaction and design are desirable is a liberating insight. It is also a matter of equity: low-fidelity prototyping allows researchers without the resources to buy an expensive robot platform, or access to a programmer, to still engage in HRI research. Greater participation means more voices and more knowledge of alternatives and how they compare with one another. Dow *et al.* note that within a collaborative team, sharing multiple designs improves exploration, group rapport, and results [14]. In the larger academic and industrial context, sharing design explorations similarly should allow researchers to explore and improve community dynamics, the general knowledge base, and eventual design outcomes [18].

Prototyping in simulation allows for studies in which all participants are remote. Remote studies make it easier for researchers to elicit responses from outside of their local area, exposing opportunities for understanding how humans in different cultures and contexts engage in and interpret interactions with robots.

Exploratory prototyping is specifically important for human-robot interaction design because the success of such interactions hinges on understanding the human side of the interaction. Wizard-of-Oz and enactment techniques can be used to draw out the many ways a human might respond to various robot behaviors, but using an existing “real robot” constrains the design space substantially. For example, the elicitation of naturalistic human typewritten and speech responses to faux intelligent natural language systems in the early 1980’s was critical to the development of actual natural language understanding in the following decades [12, 26].

2.3 When to share: how design generalizes

In the quest to obtain “validity” by requiring that the specifications for a human-in-the-loop “Wizard”-operated robot behavior be tightly specified and controlled [43], HRI researchers are likely to sacrifice generalizability. In our personal experiences as designers, we have found that design insights often generalize when we can find elements of commonality across design instances—feedback on one prototype’s affordances can offer insights into another prototype’s similar affordances. Overgeneralization, however, can result from or result in invalid assumptions and underexplored design spaces. Understanding the nature of the commonality is made possible by rapidly testing out assumptions across a broader design space.

Beyond a single project, designs can generalize across domains when there is commonality in specific areas; a chair robot prototype, for example, can offer insights into other mobile robots with similar affordances or similar motion, or other furniture, depending on the nature of the feedback and the context of use. Design processes themselves can also generalize: approaches for exploring one design space may also apply to another.

Understanding when and how design generalizes is easier when designers describe their process in substantial detail, including alternatives that were explored and discarded. Zimmerman *et al.* argue that process is important to communicate [64]—we agree, and further argue that the value lies not in process *replicability*, but in understanding how the process *context*, and its differences from the reader’s context, may influence the offered insights.

3 CASE STUDIES: EXAMPLES AND METHODS OF DESIGN EXPLORATION

The low-cost prototyping methods we describe here, and then expand on in subsequent sections, will be familiar from their frequent use in controlled studies in HRI. When used in *design explorations*, however, there are subtle shifts, beginning with goals: the intent is no longer to identify individual causal factors, but rather to cast as wide a net as possible. To be most useful to the community, we believe design explorations should clearly answer questions such as: What is the context of the exploration? What is the covered territory within the design space? What details of the exploration jumped out at the designers? What did it take to make those observations salient? Design exploration contributions are ultimately sharing the anecdotal, they are a form of “staged ethnography,” fieldwork taking place in a particular possible future [38], contrived for a specific goal.

Authors of design explorations may not always know what they are looking for, along axes that are obviously measurable: they in part seek to elucidate “unknown unknowns” [51] through their designs. As a result, in contrast to a traditional ethnography, it makes little sense to only report what the authors think might be relevant—the more inclusive of detail and peculiarity the better, with reduced emphasis on the “typical case”. A single positive example is proof of possibility.

In controlled studies, careful study designs and statistical significance testing allow scientists to infer causality and attach confidence to effects, even if the effect sizes are quite small and otherwise invisible to (often blinded) human observers or experimenters. By contrast, in design explorations and exploratory prototypes, designers are mapping out a design space by explicitly seeking large effect sizes, often without hiding their manipulations from participants, where direct observations alone can offer confidence in causality because observers can see the actual manipulation (e.g., a dancing robot) having the actual observed effect (e.g., dancing bystanders). These explorations are not necessarily intended to prove causality, but rather aim to uncover new questions to explore—whether through future design explorations or controlled studies.

3.1 Examples from HCI and HRI

A few examples from the prior HCI and HRI literature illustrate how the goals of question-uncovering and design-space-mapping are served by design explorations. Odom *et al.*’s Photobox [39] is a point design, contrived to explore the possible effects of “slow technology” in the home. From a deployment in 3 households over 14 months, the authors report that users progressed from frustration to acceptance, reflecting on past life experiences as well as on technology in the home more generally, and the exploration itself suggests future work in designing for anticipation and for better understanding what timescales trigger these feelings of frustration and acceptance. Similarly, Martelaro *et al.*’s DJ Bot [34] is a pair of point designs: Wizard of Oz-controlled speech agents for interacting with music in the car and in the home. From deployments with 5 participants, the authors report—among many other findings—that *the Wizards* make surprising use of live camera data to understand whether a song recommendation is well-received, and offer questions about methods for need-finding to future researchers.

In contrast, similar exploratory work in HRI, such as Sirkin *et al.*’s Mechanical Ottoman [49] and Takayama *et al.*’s PR2 robots “expressing thought” [53], offers design exploration couched in the language of controlled study: in both cases, anecdotally, the work is

remembered for the remarkable point designs each represents, and not for the unremarkable statistical analyses performed to meet the standards of controlled experiments. From the perspective of future designers of humanoid robots and robotic furniture, those pages might have been better used by offering more detail, for example, on the approaches that failed to elicit the desired responses or on the motivations behind specific design decisions—each of which would have been more enlightening than an inconclusive count of how many participants lifted their legs to receive an incoming ottoman.

3.2 Evaluating design explorations

In HCI, early conversations around the nature and value of design contributions to the research literature led to Zimmerman *et al.*'s identification of four criteria for evaluating “research through design” in HCI: **process**, **invention**, **relevance**, and **extensibility** [64]. The HRI community already recognizes certain contributions, such as ethnography, that are not assessed by the standards of controlled studies. Ethnographies do not exactly generalize, but still allow the community to draw inferences from several instances; similarly, a single “point design” [15]—a design illuminating one possible point in a design space—may not generalize on its own, but designers can use **abductive reasoning** to draw inferences from a limited, non-conclusive set of point designs [15].

For example, in the case studies described in subsequent sections, we explore several point designs for robotic chair interactions—involving multiple people, multiple chairs, and other furniture—that draw upon prior work with single robotic chairs. The existence of multiple point designs lets us draw inferences around methods, people, and types of interactions, such as the relative importance of timing, or degree of interactivity, compared with robot form, or even function. Sharing point designs would allow researchers to benefit from the currently-un disseminated experiences of other researchers.

Zimmerman *et al.*'s four criteria offer, in our view, a reasonable starting point for evaluating design contributions to HRI. First, the design **process** must be sufficiently well-described that other designers should be able to reproduce it, though without an expectation that reproducing the process will result in the same design outcome. To this point, we would add that the description of the process should include enough contextual detail that the reader might understand the offered insights' limitations and applications to adjacent designs *without the need for replication*. Second, contributions should be a “novel integration of various subject matters to address a specific situation” [64] (**invention**). Third, contributions should offer clear transformative potential, illuminating possible futures and articulating both how those futures differ and why we should care (**relevance**). Fourth, contributions should enable others to build on the presented work's methods or generated knowledge (**extensibility**).

3.3 Methods for design explorations in HRI

In the subsequent sections, we share case studies of three different methods we have used to explore how we want to design and program the behavior of a set of chair robots. These case studies are focused on the uses of these three exploratory prototyping methods in the service of *design explorations* rather than *conclusive results*—the intent is to provide insight into design directions worth pursuing further [19, 42].

Physical prototypes with human-in-the-loop control (§4) allow for observing responses “in the wild,” but are slow, technically challenging, and/or expensive to build. Physical prototyping enables

observable field studies whose participants are free to move about, look around, and notice and interact with things as they would in any physical environment. Creating a novel physical system can require substantial engineering work, even if it employs off-the-shelf parts. Answerable questions include “how do people react to robots in the immediate physical environment?”, “how do people respond to this motion?”, and “how do human Wizards respond to, or improvise with, common participant responses?” Follow up studies should deploy more closed-loop control to enable repeatable motion—but early on, this adds complexity and development time, and does not contribute to the initial goal of observing human behavior.

Video prototypes (§5) are pre-recorded videos of humans interacting with robots that are shown to participants, who are subsequently asked about their interpretation of the shown interaction. Answerable questions include “is the intent of this motion clear and/or interpretable?”, “what does this motion look like to a third party observer?”, “how do subtle contextual changes affect the perception of the robot's motion?”, and “how does the robot's presence, or changes in its appearance or behavior, alter how others perceive the robot as well as its interaction partners?” [47]

Virtual simulations (§6) straddle these worlds. While harder to develop than video prototypes, rendered 3D and Virtual Reality (“VR”) simulations are more immersive and allow for the observation and capture of participants' live responses to interactions. Pre-programmed motions ensure that participants are responding to precisely designed cues, while human-in-the-loop control allows observation of improvised, or better timed and adapted, Wizard responses to participant interactions. But the lack of physicality prevents participants from interacting with the objects themselves, and the simulation itself may omit real-world physical effects, such as object inertia [23] or collisions.

In the following sections, space constraints limit detail compared with what one would expect from a full exploratory prototyping contribution—these examples are illustrative of *kind*, not extent.

4 PHYSICAL PROTOTYPES WITH HUMAN-IN-THE-LOOP CONTROL

A fundamental part of human-robot interaction is understanding and designing the physical interaction that occurs between humans and the physical robot. Through our motor memory, humans have an intuitive way of understanding and engaging with physical motion in the immediate environment. Physical prototyping allows us to explore how robots' motions are perceived through this process and implicitly understood. These interactions are mediated by the *System 1* mode of thought, defined by Kahneman as an automatic or fast way of thinking, requiring no conscious effort [24].

A physical prototype should afford the kinds of physicality and interactions representative of those of the eventual robot. First, the robot should look similar to the type of object for which the motions are designed—a chair-robot should look like a chair, a sofa-robot like a sofa. A representative appearance is important because shape and size and their affordances have a large influence on how humans interpret motion. Second, the degrees of freedom of the physical prototype should match those of what the intended ultimate object would similarly be capable of.

4.1 Prior work

Low-cost physical prototyping has a rich history of use in evaluating different robotic objects, including a mechanical ottoman [48, 49],

Figure 3: A bird's-eye perspective of the cafe in which physical prototyping experiments took place.

trash barrel robot [62], chairs [27], sofa [50], drawers [37], and autonomous vehicles [44].

Similarly, the use of live enactments as a generative design method is well-established in HCI, with variations including Burns *et al.*'s "Informance Design" [6] and Odom *et al.*'s "User Enactments," in which researchers use short embodied experiences with design concepts to explore "radical alterations to technologies' roles, forms, and behaviors" [38]. As Oulasvirta *et al.* argue, a more accurate understanding of the implications of context can be supported by "bodystorming," that is, "carrying out design sessions in the original context, 'in the wild', instead of in the office." [41]

Prior work has also addressed the specific question of how to engage in rapid iteration with physical prototyping. One memorable example is Bartneck *et al.*'s exploration of how to address questions of form, purpose, social role, and environment through prototypes ranging from paper and mechanical mock-ups to Wizard of Oz remote control [3].

4.2 Case study

We used a chair robot to prototype communicative expressions at a nearby university cafeteria. The robotic chair motion was controlled by a human operator, Wizard of Oz-style, to communicate its availability to be "sat in" to cafe patrons that had just paid for their food and were looking for a place to sit (see Fig. 3). If cafe-goers noticed or interacted with the chair robot in any way, researchers would approach them after the interaction with a consent form and a survey. The survey asked about first impressions of the robot, and what was, and was not, "alright" about the interaction.

Our questions were intentionally open-ended so as not to prime for specific types of answers. Additionally, the space was video recorded with multiple cameras to allow us to reconstruct the detailed motion and interaction between the chair and participants. We found several themes in viewers' responses ($N = 33$), and many, including social engagement and safety, came up in both the "alright" and "not alright" sections of the questionnaire. Some viewers indicated their preference for more pro-social behavior with comments including "it tries to sit in the table with us, it feels like the robot knows us," while others wanted less social behavior, offering "it came to us a bit too slowly so it was surprising and a bit creepy...we weren't able to use it (we were already sitting on regular chairs)."

In our explorations, we found that it was incredibly difficult to convince participants to sit in the robot chair. From participant responses

and human controller observations, it seemed that the timing of the interactions was critically important to engaging participants.

4.3 Discussion

The most important outcome of the case study, for us, was to reveal the importance of timing in creating a successful interaction, and to suggest a critical question for follow-up exploration. Physical prototyping enables the discovery of potential interaction parameters such as timing, proxemics, and movement paths. Physical prototypes also allow people to form their judgements based on the presented affordances of both the robot (e.g. size, apparent motion characteristics) and the situation (e.g. physical and social surroundings). In general, physical prototyping is useful for developing an understanding of how humans will interact with a prototype robot physically, "in the wild," in a space populated by other real humans. Through physical prototypes, we also get close to a "representative fidelity" that allows our results to transfer to other robots and contexts more easily—but for this, we must pay special attention to designing the appearance of the robot. In the chair robot case, the form should appear as "a chair that is robotic" rather than "a robot that is a chair." This transferable design, in combination with unsuspecting bystanders—or participants, can yield a powerful research tool to explore new types of robots inhabiting various spaces.

4.3.1 Limitations. Human-in-the-loop control enables researchers to experiment with different kinds of motion without engineering automation, however, it also limits fidelity as the human operator's ability to respond might not reflect the final robot's response times or agility. These constraints also affect the researcher's ability to record and repeat identically the motions that contribute to successful interactions.

One challenge is the difficulty of controlling, recording, and then repeating the context in which all interactions take place. As a result, it can also be difficult to understand which factors in the environment—time, movement, lighting, number of bystanders—affect the interaction outcome.

5 VIDEO PROTOTYPES

Video prototypes can be very inexpensive to create, but robots depicted in the scene cannot respond to viewers of the video. A basic version of this method involves video recording a human moving a robot—little more is needed than a phone with a camera and an object that is plausibly robotic. A higher-fidelity version of this technique involves recording in front of a green screen and removing the human controller from the scene afterward, using video editing software. Removing the otherwise visible human controller imparts an illusion of autonomy to the robot, but takes more effort to create.

5.1 Prior Work

Researchers in HCI have long adapted video prototyping to their own needs. The inherent flexibility of the medium allows researchers to explore a range of visual expression and to realize ideas early in the design process [56]. Video has also been used, as we do here, in combination with Wizard of Oz techniques [33], to aid brainstorming, and as a design material itself [7].

Video prototyping as a technique in HRI has also been the subject of multiple validation studies [52, 60] showing that, although not suitable as a replacement for live field experiments, video-based

Figure 4: Top: Video prototype of an “I’m Available” interaction, using a visible human in a painter’s coverall to move the chair. Bottom: Video prototype of the “Let’s Help Out” interaction using a green screen and green suit to key out the human controller.

HRI trials do report similar results to live trials under controlled conditions [60].

5.2 Case Study

We created video prototypes of chairs communicating requests to engage with humans (such as “follow me”), or offering updates on their status (“I’m available”) using their motion alone (see Fig. 4). We used an ordinary office chair with casters, manipulated by a human, to prototype communicative expressions for our chair robot design. The videos depicted a human controller physically guiding a chair’s movement, steering from the top of its backrest, together with an actor in the same scene. In some of the videos, we digitally removed the chair’s human controller, leaving only the chair and the actor visible.

To gain an understanding of the interpretability of the chair’s motion, we recruited participants through Amazon’s Mechanical Turk to view and comment on each scene. Each set of participants was asked to watch a randomly selected set of videos, and then to answer both open-ended questions about their interpretations of the motion, such as “Describe what you saw in the video” and “How would you characterize the chair?”, and to rate, Likert-style, statements like “The chair was responsive to the person”.

Our examination of viewers’ responses yielded a few potential avenues for revision and future inquiry. We learned that certain messages were much easier to understand through motion than others, suggesting an investigation of which features of a message might yield greater or lesser interpretability. Observed correlations between ratings of chairs as *responsive* and *expressive* suggest that we might aim to design movements that capture both traits.

5.3 Discussion

Video prototypes allow researchers to evaluate different motions and scenarios with the same chair in the same physical context. Scenario-building enables quick production of 2D video, which in turn allows quick collection of interpretability information. Videos are easy to make, but do not allow for interactive discovery—they are pre-recorded, and thus the videos’ subjects can not react to observers.

The nature of video prototyping additionally provides limited opportunity for observing “in the wild” responses of bystanders to robot motion because that motion is performed by an obvious human that is only removed in post-processing.

It is possible to be flexible in several ways with video prototypes, including in the movement of the chair and its contexts, for example, in chair appearance, or number of bystanders. Further, this method is inexpensive, quick, and does not require advanced technical skills. Researchers can adjust the videos in subtle ways to find differences in perception, while keeping other elements of the video constant.

5.3.1 Limitations. Results from cognitive psychology suggest that two people participating in a conversation are far more accurate, compared with bystanders, in understanding the content of the conversation—even with the same amount of information [25]. This result can be attributed to grounding [10], a “process by which people establish a mutual understanding through their collaborative efforts to make their meanings intelligible to each other” [25]—suggesting that there are limits to what we can learn from bystanders’ understandings of an interaction from video.

Further, a short video on its own, recorded from one perspective, is not a complete representation of an interaction, and thus one needs to set the context with a cohesive story. Lastly, participants may not empathize with the people they see in the video, as they themselves are not the ones interacting with the environment.

6 VIRTUAL SIMULATIONS

Virtual simulations can be constructed in a few ways, though primarily they refer to either immersive VR worlds with human-operated or code-controlled robots, or to screen-based “video game”-style worlds, shown and interacted with using a web browser on a computer. In either case, participants are in control of some aspect of the environment, while most (or all) of the rest of the environment is simulated.

6.1 Prior Work

Simulated environments have proved a valuable tool for the development of autonomous robotic systems. High fidelity simulations have been used to train algorithms for tasks such as vehicle motion [13, 36, 46], complex object manipulation [30], and indoor navigation [28]. Interactive virtual environments have been used to train robot teleoperators [11]. There is an increasing body of work intended to make simulation for human robot interaction more accessible. Integrations between ROS and popular game engines [29] enable responsive robots to exist in the same virtual world as human participants. Leveraging work from animation has been useful in designing social robots and their behaviors [1, 2, 53].

VR environments, while immersive, also suffer from limitations that appear inherent to the form. For example, proxemics are challenging to study in VR because, although proxemic preferences exist in VR [63], those preferences appear to be inconsistent between live and VR experiences [31]. Further, walls and other physical barriers, though visible in VR, do not actually prevent motion—which can break the sense of realism.

6.2 Case Study

In a recent study, we explored how two control features—the top speed of a chair robot, and the perspective of the first-person camera view offered to the controller—affect the ways controllers navigate a crowded space to reach a goal position.

Figure 5: A: Grand Central Terminal in New York City (image courtesy Chris Bojanovich [4]). B: Grand Central simulation, rendered in WebGL. Note the chair in the foreground; the chair's destination is the white doorway along the back wall.

Using an online simulation environment, we drew on the social instincts from ($N = 30$) geographically distributed human robot controllers—online participants from Amazon Mechanical Turk—asking them to pilot a robot through a crowded 3D environment in order to elicit a variety of motion paths (see Fig. 5). To elicit socially appropriate behavior, participants were advised to “Try not to hit or scare anyone. Avoid getting in the way!” and to control the chair motion using the keyboard arrow keys. We varied the viewport’s perspective—either above and behind the chair, or at seat level facing forward—as well as chair top speed—either the same speed as the humanoid agents, or 3 times quicker.

A second set of online participants ($N = 30$) evaluated the social appropriateness of the elicited robot motions. Unlike the pilots, the evaluators were all placed in a third person perspective recreated in the same virtual world.

To develop an intuitive feel for this data, we plotted all the trajectories taken by each participant, including stops and collisions with passersby (see Fig. 6). Contrary to our expectations, we found that a high or low camera view did not seem to affect the paths controllers chose, but that controllers in the “higher top chair speed” condition collided with passersby less frequently and for shorter durations.

6.3 Discussion

The primary lesson we learned from this case study was one we did not expect: our robots were too slow, and we plan follow-on studies to explore the effect of speed on the ability of robots to avoid collisions with passersby.

The capacity to replay an interaction from an infinite variety of perspectives is both valuable and unique to this method, offering

Figure 6: Our visualization of all controllers’ trajectories, separated by affordance conditions. Trajectories are shown from above facing down, start at the left, and end at the right.

tools for prototyping, reflection and analysis of interactions that are impossible in other methods. A good simulation can provide insight into the subjective feeling of the interaction and the actual negotiation between parties; the simulated robot can respond to the participant, and respond to the participant’s responses too.

The fidelity of any simulation should be matched to the questions being asked. Low-fidelity 3D simulations can provide an opportune middle ground for design exploration between video prototyping and physical prototyping—and in some ways surpassing both: In simulation, the initial conditions of any interaction are perfectly reproducible while remaining fully interactive, enabling iterative prototyping and testing many variations with slightly different initial conditions or affordances. Being intrinsically replicable, 3D simulations also offer an avenue for testing single variable causal effects.

Like video, 3D simulation scales readily. A WebGL-based simulation can be shown to thousands of people on a crowd-work platform in hours, and simple Mobile-based VR has become increasingly accessible [32]. The capacity for scale and replication enables exploring questions more challenging for a physical prototype: simulations can offer identical interactions in different cultural contexts, and researchers can test systems in physical spaces where permission would be required and unavailable, or where robot scale prevents real-world evaluation.

6.3.1 Limitations. Though powerful, 3D simulations are not well-suited for exploring phenomena that are inherently physical or that require real-life contexts—for example, anything haptic, as we can not yet realistically simulate what it is like to touch things. Physically impossible robots are often easier to program than realistic ones, and researchers can ignore physical limitations—for better or worse. For example, while bodily harm is impossible, knowledge of that fact can encourage participants to do unrealistic things. Lastly, the constructed nature of simulations makes this method unsuitable for “field studies” of unsuspecting participants.

7 COMPARING METHODS, SHARING, TRIANGULATING

The methods described here have their own advantages and disadvantages individually, and can also complement each other in ways that provide a more complete understanding of a design space.

7.1 Comparing methods

We chose the following axes of comparison between the low-cost methods described above because they reflect, in our own judgment, the most salient differences between these techniques. Other differences bear discussion, and we hope this work will serve as a catalyst for future discussions.

7.1.1 Controlling robot motion. Video prototypes and virtual simulations provide the easiest interfaces for producing naturalistic robot motions in a variety of contexts, but at the cost of realism (due to visible human operators or physical constraints on green screen deployment) and responsiveness to participant behavior. The ease and low cost with which a variety of robot motions and interactions can be explored enables very rapid iteration at the earliest stages of design, when the mere act of viewing a video playback of your own best guess at motion provides a great deal of insight. These prototypes can be plentiful and generative, and used in the way that sketches are used elsewhere in design.

In our own work, we have often led with storyboards and video prototypes, moving to 3D simulations or physical prototyping after developing an intuition for what types of robot capabilities and affordances are worth examining in real life with real participants. The chair robot explorations described here have a longer history than the individual case studies presented, ultimately beginning with video recordings from a workshop of human-moved chairs dancing in a large atrium space.

7.1.2 Immersion and naturalistic environments. In a video prototype, compared with a simulated environment, it is easy to include real people in real physical contexts, which can help viewers imagine the real experience, and avoid the distraction and prejudice that comes with a video game-style rendering environment. Crowd-work platforms like Amazon Mechanical Turk allow researchers to easily get an understanding of participants' perspectives as bystanders to nearly any video prototype-based interaction.

In contrast, simulated environments offer the ability to observe how participants react to robot motion that responds to participant action. Simulated environments make tracking all actions and interactions easy, where a physical environment might require a complex tracking setup to collect data on participants' motions. Simulations also offer the flexibility and control to violate physical laws, put people in semi-unsafe situations (especially useful in driving studies) and program events to happen reliably and repeatably.

A physical prototype has none of these advantages, but does allow a degree of immersion impossible to otherwise achieve. A robot rapidly approaching you in the physical world is a very different experience than observing the same on a monitor or even in VR, and this can make the physical prototype the more ecologically valid choice.

Physical prototypes complement video prototypes by offering strong evidence of behavior “in the wild” with limited priming [54] and demand effects [40]. These *in-situ* studies can validate the insights generated by earlier methods and provide an understanding of how the dynamic human-robot system evolves after the first or second interactions that can be plausibly prototyped in video.

7.1.3 Technical effort. Significant engineering effort may be required to create a physical robot prototype capable of performing the same motions, with similar fidelity, that a human can perform by manipulating an armature or by moving a large, heavy “mobile” object such as a chair. Expending this effort has the most value once the potential for design insights from quicker, lower-cost methods has been exhausted.

In general, video prototypes and simulated environments are less work to develop than physical prototypes, though there are many exceptions—in cases where an existing robotics platform may be suitable, researchers may be able to trade off engineering effort for monetary expense. The cost of re-recording a video, or even modifying a simulation, is small compared with the effort required to re-engineer a physical prototype; in the earliest stages of design, iteration speed is critical.

7.2 Sharing and Triangulation

Stepping back, when we look at the three cases presented here, what is notable is how the strengths of each method help to address holes in the others, and help to buttress an overall understanding—in our cases—of how chair robots would work. Triangulating what is real is easier the more instances researchers have to draw inferences from, but often, research groups do not have the luxury of running studies on multiple platforms and featuring multiple approaches in a research paper. Ideally, the triangulation that would help the formative design process should be able to occur across groups and time. Anecdotally, some of the most valuable conversations at HRI come from sharing not-fully-formed design insights; as part of the literature, these explorations would broaden access to the insights currently only available to active researchers and other insiders.

8 CONCLUSION

In this paper, we argue for the value of exploratory prototypes as *bona fide* research contributions to HRI. We examine how rapid prototyping methods are used differently for this purpose, and argue that this demands different consideration than controlled experiments. We offer suggestions for how to consider exploratory prototype contributions, and offer a few sketches of contributions from our own work.

Further, we argue that standards for research through design that emphasize repeatability and generalizability are harmful to the ability of the HRI community to collaboratively develop new and novel ways of interacting with robots. Recognizing exploratory design lessons, even with various “faked” methods, is crucial to the design of future interactions. We believe that a “real robot” is not required for meaningful and critical contributions to the field of human-robot interaction, that this insight is critical to enabling participation from a broader and more diverse set of researchers, and that there is great value in sharing design insights that suggest interesting avenues of exploration left as future work.

ACKNOWLEDGMENTS

The authors would like to thank Ofir Yaniv and Gilie Gefen for their work on earlier drafts of this paper, as well as the many anonymous reviewers who offered feedback on earlier drafts of this work.

REFERENCES

- [1] Etienne Balit, Dominique Vaufreydaz, and Patrick Reignier. 2016. Integrating animation artists into the animation design of social robots: An open-source

- robot animation software. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*. IEEE Press, 417–418.
- [2] Etienne Balit, Dominique Vaufreydaz, and Patrick Reignier. 2018. PEAR: Prototyping Expressive Animated Robots-A framework for social robot prototyping. In *13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2018)*. VISIGRAPP.
 - [3] Christoph Bartneck and Jun Hu. 2004. Rapid prototyping for interactive robots. In *The 8th conference on intelligent autonomous systems (IAS-8)*. IOPress, 136–145.
 - [4] Chris Bojanovich. 2012. *Grand Central Terminal*. Wikipedia. https://commons.wikimedia.org/wiki/File:Grand_Central_Terminal_Lobby.jpg
 - [5] British Design Council. 2005. The 'double diamond' design process model.
 - [6] Colin Burns, Eric Dishman, William Verplank, and Bud Lassiter. 1994. Actors, hairdos & videotape—informance design. In *Conference companion on Human factors in computing systems*. ACM, 119–120.
 - [7] Jacob Buur, Thomas Binder, and Eva Brandt. 2000. Taking video beyond 'hard data' in user centred design. In *Participatory design conference*. Citeseer, 21–29.
 - [8] Bill Buxton. 2007. *Sketching User Experiences: Getting the Design Right and the Right Design (Interactive Technologies)*. Morgan Kaufmann.
 - [9] Bill Buxton. 2014. Multi-Touch Systems that I Have Known and Loved. *BillBuxton.com* (2014). <http://www.billbuxton.com/multitouchOverview.html>. Accessed: 2014-12-24.
 - [10] Herbert H Clark. 1996. *Using language*. Cambridge university press.
 - [11] Jeff Craighead, Jennifer Burke, and Robin Murphy. 2008. Using the unity game engine to develop sarge: a case study. In *Proceedings of the 2008 Simulation Workshop at the International Conference on Intelligent Robots and Systems (IROS 2008)*.
 - [12] Nils Dahlbäck, Arne Jönsson, and Lars Ahrenberg. 1993. Wizard of Oz studies—why and how. *Knowledge-based systems* 6, 4 (1993), 258–266.
 - [13] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An open urban driving simulator. *arXiv preprint arXiv:1711.03938* (2017).
 - [14] Steven Dow, Julie Fortuna, Dan Schwartz, Beth Altringer, Daniel Schwartz, and Scott Klemmer. 2011. Prototyping dynamics: sharing multiple designs improves exploration, group rapport, and results. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Acm, 2807–2816.
 - [15] Steven Dow, Wendy Ju, and Wendy Mackay. 2013. Projection, Place and Point-of-View in Research through Design. *The SAGE Handbook of Digital Technology Research* (2013), 266–285.
 - [16] Steven P Dow, Alana Glasco, Jonathan Kass, Melissa Schwarz, Daniel I Schwartz, and Scott R Klemmer. 2010. Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Transactions on Computer-Human Interaction (TOCHI)* 17, 4 (2010), 18.
 - [17] Steven P Dow, Kate Heddleston, and Scott R Klemmer. 2009. The efficacy of prototyping under time constraints. In *Proceedings of the seventh ACM conference on Creativity and cognition*. ACM, 165–174.
 - [18] Christiane Floyd. 1984. A systematic look at prototyping. In *Approaches to prototyping*. Springer, 1–18.
 - [19] TW Frick and CM Reigeluth. 1999. Formative research: A methodology for creating and improving design theories. *Instructional-design theories and models: A new paradigm of instructional theory 2* (1999), 633–652.
 - [20] Elizabeth Gerber and Maureen Carroll. 2012. The psychological experience of prototyping. *Design studies* 33, 1 (2012), 64–84.
 - [21] Anders Green, Helge Huttenrauch, and K Severinson Eklundh. 2004. Applying the Wizard-of-Oz framework to cooperative service discovery and configuration. In *RO-MAN 2004. 13th IEEE International Workshop on Robot and Human Interactive Communication (IEEE Catalog No. 04TH8759)*. IEEE, 575–580.
 - [22] Guy Hoffman and Wendy Ju. 2014. Designing robots with movement in mind. *Journal of Human-Robot Interaction* 3, 1 (2014), 89–122.
 - [23] Yuki Kado, Takanori Kamoda, Yuta Yoshiike, P Ravindra De Silva, and Michio Okada. 2010. Sociable dining table: The effectiveness of a “KonKon” interface for reciprocal adaptation. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 105–106.
 - [24] Daniel Kahneman. 2011. *Thinking, fast and slow*. Macmillan.
 - [25] Yoshihisa Kashima, Olivier Klein, and Anna E Clark. 2007. Grounding: Sharing information in social interaction. *Social communication* (2007), 27–77.
 - [26] John F Kelley. 1983. An empirical methodology for writing user-friendly natural language computer applications. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 193–196.
 - [27] Heather Knight, Timothy Lee, Brittany Hallawell, and Wendy Ju. 2017. I Get It Already! The Influence of ChairBot Motion Gestures on Bystander Response. In *Robot and Human Interactive Communication (RO-MAN), 2017 25th IEEE International Symposium on*. IEEE, 443–448.
 - [28] Eric Kolve, Roozbeh Mottaghi, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. 2017. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474* (2017).
 - [29] D Krupke, S Starke, L Einig, F Steinicke, and J Zhang. 2017. Prototyping of immersive HRI scenarios. In *International Conference on Climbing and Walking Robots and the Support Technologies for Mobile Machines*. World Scientific, 537–544.
 - [30] Youngwoon Lee, Edward S Hu, Zhengyu Yang, Alex Yin, and Joseph J Lim. 2019. IKEA Furniture Assembly Environment for Long-Horizon Complex Manipulation Tasks. *arXiv preprint arXiv:1911.07246* (2019).
 - [31] Rui Li, Marc van Almkerk, Sanne van Waveren, Elizabeth Carter, and Iolanda Leite. 2019. Comparing human-robot proxemics between virtual reality and the real world. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 431–439.
 - [32] Dan MacIsaac. 2015. Google Cardboard: A virtual reality headset for 10? *The Physics Teacher* 53, 2 (2015), 125–125.
 - [33] W. E. Mackay, D. S. Pagani, L. Faber, B. Inwood, P. Launiainen, L. Brenta, and V. Pouzol. 1995. Ariel: Augmenting paper engineering drawings. *Conference on Human Factors in Computing Systems* (1995). <http://portal.acm.org/citation.cfm?id=223763>
 - [34] Nikolas Martelaro, Sarah Mennicken, Jennifer Thom, Henriette Cramer, and Wendy Ju. 2020. Using Remote Controlled Speech Agents to Explore Music Experience in Context. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 2065–2076.
 - [35] Elias Matsas, George-Christopher Vosniakos, and Dimitris Batras. 2018. Prototyping proactive and adaptive techniques for human-robot collaboration in manufacturing using virtual reality. *Robotics and Computer-Integrated Manufacturing* 50 (2018), 168–180.
 - [36] Wei Meng, Yuchao Hu, Jiaxin Lin, Feng Lin, and Rodney Teo. 2015. ROS+ unity: An efficient high-fidelity 3D multi-UAV navigation and control simulator in GPS-denied environments. In *IECON 2015-41st Annual Conference of the IEEE Industrial Electronics Society*. IEEE, 002562–002567.
 - [37] Brian Mok, Stephen Yang, David Sirkin, and Wendy Ju. 2014. Empathy: interactions with emotive robotic drawers. In *2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 250–251.
 - [38] William Odom, John Zimmerman, Scott Davidoff, Jodi Forlizzi, Anind K Dey, and Min Kyung Lee. 2012. A fieldwork of the future with user enactments. In *Proceedings of the Designing Interactive Systems Conference*. ACM, 338–347.
 - [39] William T Odom, Abigail J Sellen, Richard Banks, David S Kirk, Tim Regan, Mark Selby, Jodi L Forlizzi, and John Zimmerman. 2014. Designing for slowness, anticipation and re-visitation: a long term field study of the photobox. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1961–1970.
 - [40] Martin T Orne. 1962. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American psychologist* 17, 11 (1962), 776.
 - [41] Antti Oulasvirta, Esko Kurvinen, and Tomi Kankainen. 2003. Understanding contexts by being there: case studies in bodystorming. *Personal and ubiquitous computing* 7, 2 (2003), 125–134.
 - [42] Edward L Palmer. 1973. Formative research in the production of television for children. (1973).
 - [43] Laurel D Riek. 2012. Wizard of oz studies in hri: a systematic review and new reporting guidelines. *Journal of Human-Robot Interaction* 1, 1 (2012), 119–136.
 - [44] Dirk Rothenbücher, Jany Li, David Sirkin, Brian Mok, and Wendy Ju. 2016. Ghost driver: A field study investigating the interaction between pedestrians and driverless vehicles. In *Robot and Human Interactive Communication (RO-MAN), 2016 25th IEEE International Symposium on*. IEEE, 795–802.
 - [45] Michael Schrage. 1999. *Serious play: How the world's best companies simulate to innovate*. Harvard Business Press.
 - [46] Shital Shah, Debadepta Dey, Chris Lovett, and Ashish Kapoor. 2018. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and service robotics*. Springer, 621–635.
 - [47] David Sirkin and Wendy Ju. 2012. Consistency in physical and on-screen action improves perceptions of telepresence robots. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. 57–64.
 - [48] David Sirkin and Wendy Ju. 2014. Using embodied design improvisation as a design research tool. In *Proceedings of the international conference on Human Behavior in Design (HbID 2014)*, Ascona, Switzerland.
 - [49] David Sirkin, Brian Mok, Stephen Yang, and Wendy Ju. 2015. Mechanical ottoman: how robotic furniture offers and withdraws support. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 11–18.
 - [50] Marco Spadafora, Victor Chahuneau, Nikolas Martelaro, David Sirkin, and Wendy Ju. 2016. Designing the behavior of interactive objects. In *Proceedings of the TEI'16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 70–77.
 - [51] Alistair Sutcliffe and Pete Sawyer. 2013. Requirements elicitation: Towards the unknown unknowns. In *2013 21st IEEE International Requirements Engineering Conference (RE)*. IEEE, 92–104.
 - [52] Dag Sverre Syrdal, Nuno Otero, and Kerstin Dautenhahn. 2008. Video prototyping in human-robot interaction: Results from a qualitative study. (2008).
 - [53] Leila Takayama, Doug Dooley, and Wendy Ju. 2011. Expressing thought: improving robot readability with animation principles. In *Human-Robot Interaction (HRI), 2011 6th ACM/IEEE International Conference on*. IEEE, 69–76.
 - [54] Steven P Tipper. 1985. The negative priming effect: Inhibitory priming by ignored objects. *The quarterly journal of experimental psychology* 37, 4 (1985), 571–590.
 - [55] Maryam Tohid, William Buxton, Ronald Baecker, and Abigail Sellen. 2006. Getting the right design and the design right. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 1243–1252.
 - [56] Laurie Vertelney. 1989. Using video to prototype user interfaces. *ACM SIGCHI Bulletin* 21, 2 (1989), 57–61.
 - [57] Matthew B Wall, Karl T Ulrich, and Woodie C Flowers. 1992. Evaluating prototyping technologies for product design. *Research in Engineering Design* 3, 3 (1992), 163–177.
 - [58] Astrid Weiss, Regina Bernhaupt, Daniel Schwaiger, Martin Altmaninger, Roland Buchner, and Manfred Tscheligi. 2009. User experience evaluation with a

- Wizard of Oz approach: Technical and methodological considerations. In *2009 9th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 303–308.
- [59] Jacob O Wobbrock and Julie A Kientz. 2016. Research contributions in human-computer interaction. *interactions* 23, 3 (2016), 38–44.
- [60] Sarah Woods, Michael Walters, Kheng Lee Koay, and Kerstin Dautenhahn. 2006. Comparing human robot interaction scenarios using live and video based methods: towards a novel methodological approach. In *9th IEEE International Workshop on Advanced Motion Control, 2006*. IEEE, 750–755.
- [61] Sarah N Woods, Michael L Walters, Kheng Lee Koay, and Kerstin Dautenhahn. 2006. Methodological issues in HRI: A comparison of live and video-based methods in robot to human approach direction trials. In *Robot and Human Interactive Communication. ROMAN*. IEEE, 51–58.
- [62] Stephen Yang, Brian Ka-Jun Mok, David Sirkin, Hillary Page Ive, Rohan Maheshwari, Kerstin Fischer, and Wendy Ju. 2015. Experiences developing socially acceptable interactions for a robotic trash barrel. In *Robot and Human Interactive Communication (RO-MAN), 2015 24th IEEE International Symposium on*. IEEE, 277–284.
- [63] Nick Yee, Jeremy N Bailenson, Mark Urbanek, Francis Chang, and Dan Merget. 2007. The unbearable likeness of being digital: The persistence of nonverbal social norms in online virtual environments. *CyberPsychology & Behavior* 10, 1 (2007), 115–121.
- [64] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 493–502.